

Arvados - Bug #18486

Docker containers are always removed

11/29/2021 02:02 PM - Tom Schoonjans

Status: New	Start date:
Priority: Normal	Due date:
Assigned To:	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version:	
Description Observed in Arvados 2.3.1: When trying to debug a CWL workflow running on the Docker container runtime, it appears that the Docker containers are automatically removed after they have finished running. This happens regardless of the arvados-docker-cleaner service running or the RemoveStoppedContainers setting in its config file.	
Related issues: Related to Arvados - Feature #12900: [Crunch2] [crunch-run] Prune old images ... New	

History

#1 - 12/06/2021 09:01 PM - Peter Amstutz

- Release deleted (45)

As a process note, we use the "Release" field to designate which release a bug is being fixed in, not as the release the bug was found.

#2 - 12/06/2021 09:05 PM - Peter Amstutz

You want the (stopped) containers themselves to stick around, not just the images? In general we avoid that because you can fill up your scratch space very quickly, and users typically don't have access to compute nodes with containers anyway.

However we could add some kind of admin-level configuration option for debugging in those cases where the users *do* have access to the compute node.

You might also be interested in the container shell access feature:

<https://doc.arvados.org/v2.3/install/container-shell-access.html>

<https://doc.arvados.org/v2.3/user/debugging/container-shell-access.html>

#3 - 12/07/2021 09:32 AM - Tom Schoonjans

Peter Amstutz wrote:

You want the (stopped) containers themselves to stick around, not just the images? In general we avoid that because you can fill up your scratch space very quickly, and users typically don't have access to compute nodes with containers anyway.

However we could add some kind of admin-level configuration option for debugging in those cases where the users *do* have access to the compute node.

You might also be interested in the container shell access feature:

<https://doc.arvados.org/v2.3/install/container-shell-access.html>

<https://doc.arvados.org/v2.3/user/debugging/container-shell-access.html>

Yes when we ran into trouble with the Singularity runtime last week, I gave the Docker runtime a try instead, but couldn't debug any issues as the containers were removed immediately after they finished running. This seems to contradict the note in <https://doc.arvados.org/v2.3/install/crunch2/install-compute-node-docker.html#docker-cleaner>, which states that the arvados-docker-cleaner daemon is responsible for cleaning up Docker containers (and images), meaning that no containers should ever get removed if the daemon is not running, or when "RemoveStoppedContainers": "never" is added to its config file.

Not really an issue for us, since we got the Singularity runtime up and running again after Tom's analysis of the problem and suggested fix, but thought it would be good for you to know about this.

#4 - 12/07/2021 03:31 PM - Peter Amstutz

When the container stops, we call ContainerRemove(). That's by design.

The docker-cleaner service is a bit of a legacy. I think you're right that the documentation is a little misleading. It's been our intention to get rid of docker-cleaner entirely and have crunch-run be responsible for cleaning up containers and container images when it starts ([#12900](#)). That would be closer to what you want, then the container would stick around at least until the next container starts (or you could drain the node to prevent new jobs from being scheduled).

On the other hand, singularity doesn't leave containers or container images around at all after it stops, and we load the singularity image from arv-mount, so none of this applies to the singularity case.

#5 - 12/07/2021 03:31 PM - Peter Amstutz

- Related to Feature #12900: [Crunch2] [crunch-run] Prune old images before installing image for current container added